

Configuring Hardware to Run ArcIMS: How Much is Enough?

Tony Morse and Ben Britton
Idaho Department of Water Resources
1301 North Orchard St.
Boise, Idaho 83706

Abstract

The Idaho Department of Water Resources has begun to serve ArcIMS applications for itself and for other organizations. The budgeting constraints of state government mean that any hardware designed to support ArcIMS needs to be robust enough to function adequately for several years. In assessing the department's possible ArcIMS usage, IDWR personnel concluded that additional hardware would be needed, and that the hardware would have to be scalable. The hardware configuration planned for supporting ArcIMS consists of a large server and a separate, large workstation-class PC running NT Server and all the ArcIMS software except the servlet connector. The Hewlett Packard LH6000 Server will run SDE and SQL Server, and serve data. The ArcIMS PC will have the fastest available CPU and will run only ArcIMS. It will function as the Application Server and the Spatial Server. It will store and serve images through ArcIMS until the release of SDE 8.1. The server and the ArcIMS PC will be connected by a dedicated, isolated 100 Mb network. A fully-configured LH 6000 has 4 CPUs, 8 GB of RAM, and 436 GB of disk. IDWR will buy it with 4 CPUs, 8 GB of RAM and all 436 GB of disk. If more CPUs are needed, they will be added at a later date. The ArcIMS PC will have a single P4/1.5 GHz CPU, 1 GB of RAM, and two 36.4 GB SCSI disks. As the IMS workload increases, Moore's law will allow this component to be replaced in a little over a year with a CPU that runs twice as fast. Tests run at IDWR using existing hardware showed that a split-server configuration will run ArcIMS up to 300% faster than a one-server configuration for essentially the same price.

Introduction

The Idaho Department of Water Resources (IDWR) has begun serving ArcIMS 3.0 applications from its GIS server. The GIS server is a Hewlett Packard LH3 with two PIII-500 MHz CPUs, 512 MB of RAM, and twelve 18.2 GB SCSI disks in RAID 5. The hardware runs ArcIMS, SQL Server, and the Spatial Database Engine (SDE), and also is a departmental file-server for 100 GB of vector and raster GIS data.

Presently, ArcIMS performance is sluggish and unstable on the LH3, which can support more RAM, but is otherwise at its limit of scalability, and will not be able to successfully support the present half-dozen ArcIMS applications, let alone future applications already being requested by IDWR business-process managers. Therefore, IDWR GIS personnel undertook an effort to determine the hardware that would be needed to run ArcIMS with a reasonable response time and be scalable to accommodate future application development.

Several options for configuring hardware to run ArcIMS were considered. The best configuration is a two-server ArcIMS implementation consisting of a powerful, single-purpose workstation-class computer that runs only ArcIMS and is connected by a dedicated, isolated 100 Mb/S line to a robust, scalable, server running SDE and SQL server. This conclusion is based on four factors: 1) conversations with ESRI personnel on server configuration and sizing; 2) tests run on

IDWR's existing configuration and on a proposed two-computer configuration; 3) the need for scalability of the ArcIMS configuration; and 4) cost.

Conversations with ESRI Personnel

Conversations with ESRI personnel were with both Redlands and the Olympia office, and covered our existing configuration, our data volume and type, and our ArcIMS needs. Redlands personnel derive ideal sizing factors from empirical testing. The ESRI-Olympia personnel have run their own tests and helped various Northwest customers configure ArcIMS sites. In general, the server sizes recommended by Redlands were smaller than those recommended by the Olympia office, perhaps due to the testing circumstances at Redlands being more ideal than the Olympia "real-world" circumstances. Redlands personnel suggested the two-server configuration described here.

Fundamentally, the only disagreement between Redlands and Olympia was over the minimum size needed for the hardware. However, since IDWR needs a robust ArcIMS configuration that will carry the load for at least three years, the minimum configuration is not of interest. The important considerations are that architecture will support our needs, and provide scalability and cost-effective performance.

Test Results

As part of our effort to determine the best hardware configuration for running ArcIMS, IDWR GIS personnel ran two tests, which are summarized below. The present IDWR configuration, which is diagramed in Figure 2, has ArcIMS, SQL Server and SDE on the GIS data server. ESRI personnel recommended separating ArcIMS and SDE on different servers to increase ArcIMS response time. To test this recommendation, ArcIMS was split off from the main data-server. Microsoft's Personal Web Server and ArcIMS were installed on a workstation and the ArcIMS response was benchmarked for two configurations: 1) the existing configuration with ArcIMS on the server with SDE and SQL, and 2) the new, two-computer configuration in which ArcIMS is isolated on a separate platform from SDE/SQL Server. The ArcIMS workstation for the test is an 800 MHz, Pentium III with 512 MB of memory, running NT 4.0 Workstation. It was connected to the data server through the IDWR 100Mb/S LAN.

Two datasets were tested with each configuration. The server-log excerpts show the times necessary to provide the maps for those tests. The logs show the time necessary to render a map when drawn at the initial map extent.

The first test used data illustrated by Figure 1, which shows a State-wide Landsat mosaic and shape files of IDWR administrative basins and major cities, all of which were read from A10 over the network. There are no SDE layers presented at this zoom level. Table 1 shows the two server log-files from this test. For our existing configuration (ArcIMS and SDE/SQL Server on the same server) the time to render was 5.06 seconds. For the split configuration (ArcIMS and SDE/SQL Server on separate servers), the time to render was 2.5 seconds, just over 100% faster. This speed increase is startling when two factors are considered: 1) the workstation's 800 MHz CPU is 20% slower than the two 500 MHz CPUs (ArcIMS runs multi-threaded), and 2) the workstation is retrieving all the data over our network, while the server has all the data on local RAID disks.

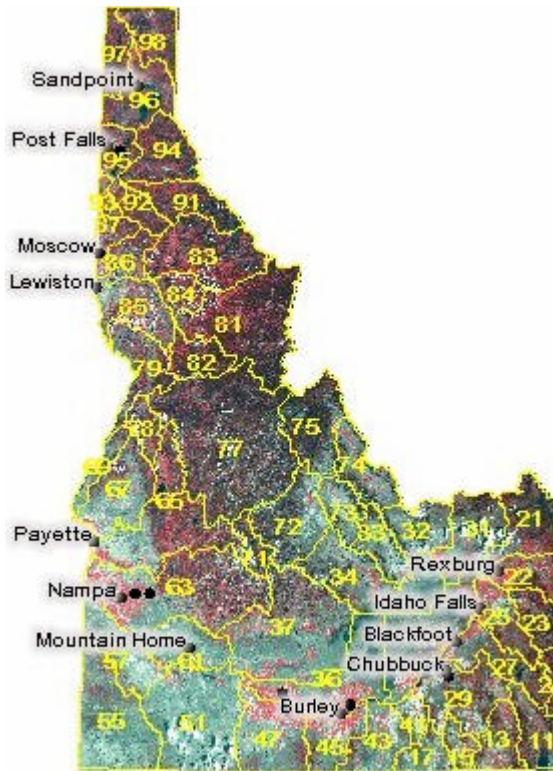


Figure 1. the first data set tested: a State-wide Landsat mosaic and shape files of IDWR administrative basins and major cities.

<p>from ImageServer_DWR509_332.log</p> <pre> <GET_IMAGE> [Tue Feb 13 11:32:15 2001][332 313 INFO1] SERVICE: Adj [Tue Feb 13 11:32:15 2001][332 313 INFO2] AXL Parse Time: 0.000000s [Tue Feb 13 11:32:15 2001][332 313 INFO2] RENDERER SETUP: 0.010000s [Tue Feb 13 11:32:15 2001][332 313 INFO2] IMAGE LAYER: Landsat Image [Tue Feb 13 11:32:15 2001][332 313 INFO2] IMAGE SEARCH TIME: 0.000000s [Tue Feb 13 11:32:17 2001][332 313 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 11:32:17 2001][332 313 INFO2] IMAGE RETRIEVAL TIME: 1.842000s [Tue Feb 13 11:32:17 2001][332 313 INFO2] FEATURE LAYER: Administrative Basins [Tue Feb 13 11:32:17 2001][332 313 INFO2] DATA SEARCH TIME: 0.010000s [Tue Feb 13 11:32:18 2001][332 313 INFO2] GR FEATURES PROCESSED: 53 [Tue Feb 13 11:32:18 2001][332 313 INFO2] DATA RETRIEVAL TIME: 0.511000s [Tue Feb 13 11:32:18 2001][332 313 INFO2] FEATURE LAYER: Major Cities [Tue Feb 13 11:32:18 2001][332 313 INFO2] DATA SEARCH TIME: 0.010000s [Tue Feb 13 11:32:18 2001][332 313 INFO2] GR FEATURES PROCESSED: 20 [Tue Feb 13 11:32:18 2001][332 313 INFO2] DATA RETRIEVAL TIME: 0.010000s [Tue Feb 13 11:32:18 2001][332 313 INFO2] TOTAL PROCESSING TIME: 2.413000s [Tue Feb 13 11:32:18 2001][332 313 INFO2] OUTPUT TIME: 0.070000s [Tue Feb 13 11:32:18 2001][332 313 INFO3] RESPONSE: <?xml version="1.0"?> <ARCCXML version="1.0"> <RESPONSE> <IMAGE> <ENVELOPE minx="-311222.706524262380000" miny="98702.492188000004000" maxx="1296915.542983262300000" maxy="882311.842775000030000" /> <OUTPUT file="D:\ArcIMS\output\Adj_DWR5093323135.jpg" url="http://dwr509/mapoutput/Adj_DWR5093323135.jpg" /> </IMAGE> </RESPONSE> </ARCCXML> [Tue Feb 13 11:32:18 2001][332 313 INFO2] Total Request Time: 2.503000s [Tue Feb 13 11:32:18 2001][332 313 INFO1] End Request </pre>	<p>from ImageServer_A10_142.log</p> <pre> <GET_IMAGE> [Tue Feb 13 11:34:19 2001][142 242 INFO1] SERVICE: Adj1 [Tue Feb 13 11:34:19 2001][142 242 INFO2] AXL Parse Time: 0.485000s [Tue Feb 13 11:34:19 2001][142 242 INFO2] RENDERER SETUP: 0.016000s [Tue Feb 13 11:34:19 2001][142 242 INFO2] IMAGE LAYER: Landsat Image [Tue Feb 13 11:34:19 2001][142 242 INFO2] IMAGE SEARCH TIME: 0.109000s [Tue Feb 13 11:34:21 2001][142 242 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 11:34:21 2001][142 242 INFO2] IMAGE RETRIEVAL TIME: 2.297000s [Tue Feb 13 11:34:21 2001][142 242 INFO2] FEATURE LAYER: Administrative Basins [Tue Feb 13 11:34:22 2001][142 242 INFO2] DATA SEARCH TIME: 0.391000s [Tue Feb 13 11:34:23 2001][142 242 INFO2] GR FEATURES PROCESSED: 53 [Tue Feb 13 11:34:23 2001][142 242 INFO2] DATA RETRIEVAL TIME: 0.921000s [Tue Feb 13 11:34:23 2001][142 242 INFO2] FEATURE LAYER: Major Cities [Tue Feb 13 11:34:23 2001][142 242 INFO2] DATA SEARCH TIME: 0.000000s [Tue Feb 13 11:34:23 2001][142 242 INFO2] GR FEATURES PROCESSED: 20 [Tue Feb 13 11:34:23 2001][142 242 INFO2] DATA RETRIEVAL TIME: 0.016000s [Tue Feb 13 11:34:23 2001][142 242 INFO2] TOTAL PROCESSING TIME: 4.375000s [Tue Feb 13 11:34:23 2001][142 242 INFO2] OUTPUT TIME: 0.157000s [Tue Feb 13 11:34:23 2001][142 242 INFO3] RESPONSE: <?xml version="1.0"?> <ARCCXML version="1.0"> <RESPONSE> <IMAGE> <ENVELOPE minx="-89083.313975859259000" miny="98702.492188000004000" maxx="1074776.150434859300000" maxy="882311.842775000030000" /> <OUTPUT file="c:\arcims\output\Adj1_A10142242346.jpg" url="http://a10/Mapoutput/Adj1_A10142242346.jpg" /> </IMAGE> </RESPONSE> </ARCCXML> [Tue Feb 13 11:34:23 2001][142 242 INFO2] Total Request Time: 5.063000s [Tue Feb 13 11:34:23 2001][142 242 INFO1] End Request </pre>
---	---

Table 1. Log files comparing the two configurations for a vector-on raster map.

The second test involved drawing points of diversion on a MrSID file of scanned aerial photography. The logs showing the time necessary to render this image are shown on the next page. Figure 2 shows the resulting ArcIMS map



Figure 2. The rendered vector-on-raster image from the second data set.

<pre>from ImageServer_DWR509_332.log </GET_IMAGE> [Tue Feb 13 12:02:37 2001][332 314 INFO1] SERVICE: Adj [Tue Feb 13 12:02:37 2001][332 314 INFO2] AXL Parse Time: 0.000000s [Tue Feb 13 12:02:37 2001][332 314 INFO2] RENDERER SETUP: 0.010000s [Tue Feb 13 12:02:37 2001][332 314 INFO2] IMAGE LAYER: Landsat Image [Tue Feb 13 12:02:37 2001][332 314 INFO2] IMAGE SEARCH TIME: 0.000000s [Tue Feb 13 12:02:37 2001][332 314 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 12:02:37 2001][332 314 INFO2] IMAGE RETRIEVAL TIME: 0.121000s [Tue Feb 13 12:02:37 2001][332 314 INFO2] IMAGE LAYER: Adjudication Imagery [Tue Feb 13 12:02:37 2001][332 314 INFO2] IMAGE SEARCH TIME: 0.000000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 12:02:39 2001][332 314 INFO2] IMAGE RETRIEVAL TIME: 1.422000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] FEATURE LAYER: Administrative Basins [Tue Feb 13 12:02:39 2001][332 314 INFO2] DATA SEARCH TIME: 0.010000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] GR FEATURES PROCESSED: 3 [Tue Feb 13 12:02:39 2001][332 314 INFO2] DATA RETRIEVAL TIME: 0.030000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] FEATURE LAYER: Cities [Tue Feb 13 12:02:39 2001][332 314 INFO2] DATA SEARCH TIME: 0.020000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] GR FEATURES PROCESSED: 3 [Tue Feb 13 12:02:39 2001][332 314 INFO2] DATA RETRIEVAL TIME: 0.000000s [Tue Feb 13 12:02:39 2001][332 314 INFO2] FEATURE LAYER: Point of Diversion - Recommended [Tue Feb 13 12:02:41 2001][332 314 INFO2] DATA SEARCH TIME: 2.183000s [Tue Feb 13 12:02:42 2001][332 314 INFO2] GR FEATURES PROCESSED: 2060 [Tue Feb 13 12:02:42 2001][332 314 INFO2] DATA RETRIEVAL TIME: 1.422000s [Tue Feb 13 12:02:42 2001][332 314 INFO2] TOTAL PROCESSING TIME: 5.208000s [Tue Feb 13 12:02:42 2001][332 314 INFO2] OUTPUT TIME: 0.070000s [Tue Feb 13 12:02:42 2001][332 314 INFO3] RESPONSE: <?xml version="1.0"?> <ARCXML version="1.0"> <RESPONSE> <IMAGE> <ENVELOPE minx="433535.396597139010000" miny="169230.799325984500000" maxx="475035.394231709360000" maxy="197172.133473447930000" /> <OUTPUT file="D:\ArcIMS\output\Adj_DWR50933231412.jpg" url="http://dwr509/mapoutput/Adj_DWR50933231412.jpg" /> </IMAGE> </RESPONSE> </ARCXML> [Tue Feb 13 12:02:42 2001][332 314 INFO2] Total Request Time: 5.298s [Tue Feb 13 12:02:42 2001][332 314 INFO1] End Request</pre>	<pre>from ImageServer_A10_142.log </GET_IMAGE> [Tue Feb 13 12:00:13 2001][142 235 INFO1] SERVICE: Adj1 [Tue Feb 13 12:00:13 2001][142 235 INFO2] AXL Parse Time: 0.032000s [Tue Feb 13 12:00:13 2001][142 235 INFO2] RENDERER SETUP: 0.000000s [Tue Feb 13 12:00:13 2001][142 235 INFO2] IMAGE LAYER: Landsat Image [Tue Feb 13 12:00:13 2001][142 235 INFO2] IMAGE SEARCH TIME: 0.000000s [Tue Feb 13 12:00:13 2001][142 235 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 12:00:13 2001][142 235 INFO2] IMAGE RETRIEVAL TIME: 0.093000s [Tue Feb 13 12:00:13 2001][142 235 INFO2] IMAGE LAYER: Adjudication Imagery [Tue Feb 13 12:00:13 2001][142 235 INFO2] IMAGE SEARCH TIME: 0.000000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] BLOCKS RETRIEVED: 1 [Tue Feb 13 12:00:15 2001][142 235 INFO2] IMAGE RETRIEVAL TIME: 2.594000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] FEATURE LAYER: Administrative Basins [Tue Feb 13 12:00:15 2001][142 235 INFO2] DATA SEARCH TIME: 0.016000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] GR FEATURES PROCESSED: 3 [Tue Feb 13 12:00:15 2001][142 235 INFO2] DATA RETRIEVAL TIME: 0.031000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] FEATURE LAYER: Cities [Tue Feb 13 12:00:15 2001][142 235 INFO2] DATA SEARCH TIME: 0.016000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] GR FEATURES PROCESSED: 3 [Tue Feb 13 12:00:15 2001][142 235 INFO2] DATA RETRIEVAL TIME: 0.000000s [Tue Feb 13 12:00:15 2001][142 235 INFO2] FEATURE LAYER: Point of Diversion - Recommended [Tue Feb 13 12:00:22 2001][142 235 INFO2] DATA SEARCH TIME: 6.656000s [Tue Feb 13 12:00:29 2001][142 235 INFO2] GR FEATURES PROCESSED: 2021 [Tue Feb 13 12:00:29 2001][142 235 INFO2] DATA RETRIEVAL TIME: 6.687000s [Tue Feb 13 12:00:29 2001][142 235 INFO2] TOTAL PROCESSING TIME: 16.109000s [Tue Feb 13 12:00:29 2001][142 235 INFO2] OUTPUT TIME: 0.141000s [Tue Feb 13 12:00:29 2001][142 235 INFO3] RESPONSE: <?xml version="1.0"?> <ARCXML version="1.0"> <RESPONSE> <IMAGE> <ENVELOPE minx="433970.106650954690000" miny="167774.331608518780000" maxx="475281.361213409520000" maxy="195588.587839196840000" /> <OUTPUT file="c:\arcims\output\Adj1_A10142235393.jpg" url="http://a10/Mapoutput/Adj1_A10142235393.jpg" /> </IMAGE> </RESPONSE> </ARCXML> [Tue Feb 13 12:00:29 2001][142 235 INFO2] Total Request Time: 16.282s [Tue Feb 13 12:00:29 2001][142 235 INFO1] End Request</pre>
---	--

Table 2. The log files for rendering vectors on a MrSID image.

In this test, the speed difference is dramatic. The existing configuration took 16.28 seconds versus 5.29 for the split configuration, just over a 300% increase. Again, -- the workstation's 800 MHz CPU is 20% slower than the two 500 MHz CPUs, and the workstation is retrieving all the data over our network, this time retrieving SDE features, while the server has all the data on local RAID disks.

The tests, while not exhaustive, are convincing: the split configuration is much faster than the single-server configuration. Nevertheless, the tests address only one-half of the issue. The other half is the absolute rendering speed needed to run an ArcIMS installation serving multiple users. ESRI personnel recommended sizing the ArcIMS server to achieve a rendering time of 2 seconds or less, a goal that the 800 MHz workstation failed to meet.

The split configuration can run faster than the PIII-800 by isolating the ArcIMS-SDE communications link, and by configuring the ArcIMS server with the fastest available CPU and at least 1 GB of RAM. The ArcIMS computer and the SDE/SQL Server computer need to communicate over a dedicated, isolated 100 Kb/s line rather than over the IDWR network, which is straight-forward to implement.

The results of the testing confirm the recommendation of ESRI personnel about the recommended configuration of ArcIMS, and provide a quantitative measure of the performance increase we can expect over our existing configuration.

Need for Scalability

The system needs to be as robust as possible and as scalable as possible if it is to function well for at least three years. Splitting ArcIMS from the SDE and SQL Server platform serves this need for scalability. The SDE/SQL Server platform will be an HP LH6000. The LH6000 can take up to 6 CPUs and 8 GB of RAM. Based on our experience with our existing GIS server, which is a 2-CPU server with 512 MB RAM, the LH6000 with 4 CPUs and 8 GB of RAM will perform well. Our present GIS server now is overloaded running ArcIMS, SDE, SQL Server, and providing GIS file services. If divide the load by keeping the file services on the present server, putting SDE and SQL Server on the LH6000, and putting ArcIMS on a large workstation-class server, we will have a division and specialization of labor among the computers. The present GIS server will no longer be over-burdened and will perform well in handling file services and in processing relatively small data sets as a development server. The LH6000 will handle only SDE and SQL Server, and will be scalable. Because of Moore's Law, we will be able completely replace the ArcIMS server as often as every 15 months with a new workstation that will be twice as fast at a price-point of about \$5,000; less if the SCSI disk sub-system is cannibalized from the present computer.

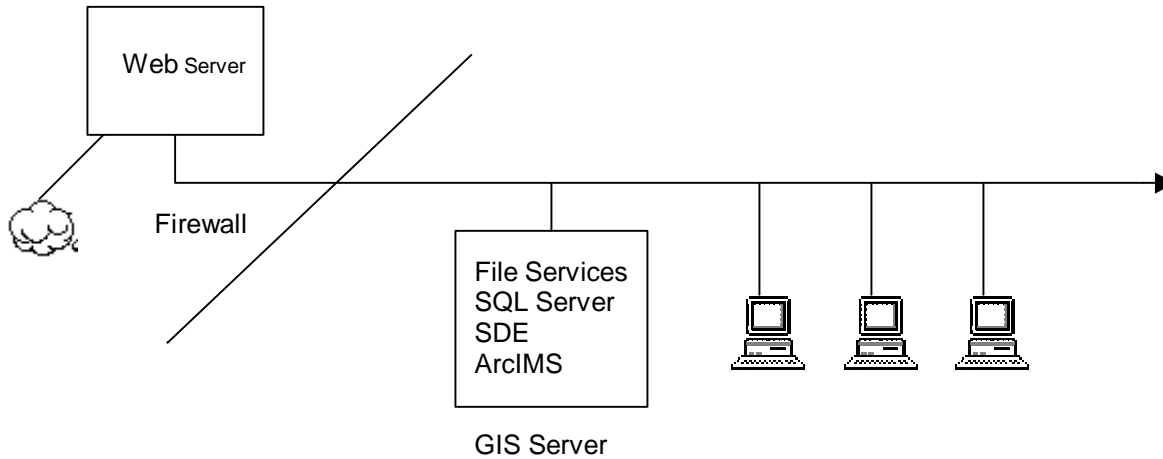
Cost

The cost for software licensing makes the separate ArcIMS workstation cost-effective. ArcIMS is licensed by the CPU at \$7,500 for one CPU and \$5,000 for each successive CPU. This means that running on a single CPU computer is \$7,500 and running ArcIMS on a quad-CPU computer costs \$22,500, just for the ArcIMS licensing.

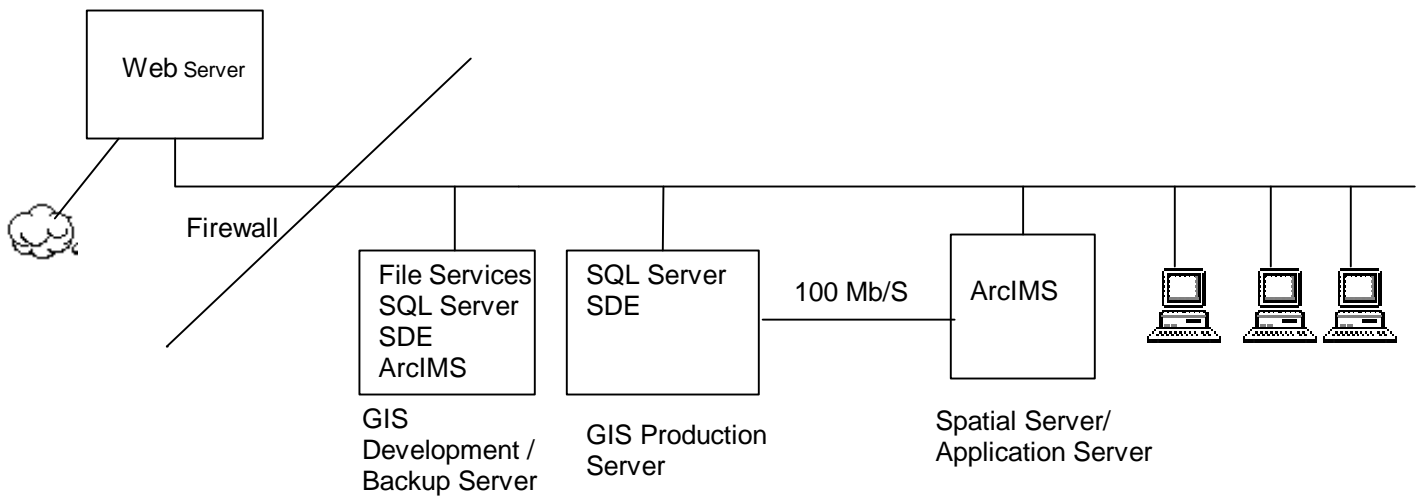
From the testing, it is clear that ArcIMS runs on a single, 800 MHz-CPU workstation two to three times faster than it will on a dual, 500 MHz server. A Dell Precision Workstation 330 with a 1,500 MHz CPU, 1024 MB RAM, and two 36.4 GB SCSI disks costs about \$5,100, which makes running ArcIMS on a separate workstation-class server approximately \$10,000 less expensive than it is to run on a quad-processor server. The \$10,000-saving alone makes the configuration worthwhile.

Specific Hardware Configuration

Present Configuration



Proposed Configuration



The specific hardware proposed for the GIS Production Server

Hewlett Packard LH6000 Server - \$43,300 (Computerland quote dated 2/7/01)

LH6000 Netserver SE
4 Xeon 700 MHz 1MB CPUs
8 GB 133 MHz DIMM RAM
Two 10/100TX PCI LAN card2
NetRAID 128 MB Controller card
NetRAID 4M Manual kit
12 36.4 GB 10K ULTRA 3 WD SCSI3 disks
Internal drive cage upgrade
Windows 2000 Server

The specific hardware proposed for the ArcIMS Server

Dell Precision Workstation 330 - approximately \$5,100

P4/1.5 GHz
1 GB PC 800 ECC RDRAM
Nvidia tnt Pro
Ultra 160/M SCSI Controller
Two 36.4 GB drive (10,000 rpm)
No monitor
Windows 2000 Server
Two 32/64-bit Copper gigabit PCI network cards
3-year parts and labor
Windows 2000 Workstation (removed)
Windows 2000 Server (additional cost)

The existing GIS Development/Backup Server

Hewlett Packard LH 3 Server
Two PIII/500 MHz CPUs
500 MB RAM
Twelve 18.4 GB SCSI 3 Disks
Windows NT 4.0 Server

Postscript

IDWR purchased the two computers described above. IDWR personnel encountered a significant problem with ServletExec on the Dell. Each time ServletExec was installed, it halted the world-wide-web publishing services within Microsoft IIS on the web server. The effort returned a Dr. Watson error as follows: *"An application error has occurred and an application error log is being created. inetinfo.exe Exception: access violation (0xc0000005), Address: 0x500bf974"*. That message was followed by : *"Could not start World Wide Web Publishing Services on \\A018. Error 1067: The process terminated unexpectedly."*

The error was returned first under Windows 2000 SP1. Uninstalling Windows 2000 and installing NT 4.0 Server, SP6a made no difference; the error persisted. Switching-out the Dell P4-1500 to using instead a Micron P3-800 solved the problem immediately.

ServletExec runs in the Java RunTime Environment. ServletExec initiates the Java console (javaw.exe) and that executable causes the access violation. We have tested ServletExec 3.0 and 3.1, in combination with JRE 1.1.8, 1.2.2 and 1.3, with the same result. Since the problem was solved by running the software on a Pentium III, the Pentium 4-architecture is suspect. In any case, a solution has not yet been found.